SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Bridging Neural Perception And Symbolic Reasoning For Intelligent Industrial Process Automation

N. Nivetha[1*], M. Gayathri[2], Najmitdinov Akhadkhon Khamitdkhanovich[3], Dr. Lalit Sachdeva[4]

[1*]Assistant Professor, Department of Computer Science, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Chennai, Tamil Nadu, India. E-mail: nivethan@maher.ac.in

[2]Assistant Professor, Department of Management Studies, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Chennai, Tamil Nadu, India. E-mail: gayathrimba@maher.ac.in

[3]Rector, Faculty of Business Administration, Business Administration, Turan International University, Namangan, Uzbekistan. E-mail: a.najmitdinov@turan-edu.kz, https://orcid.org/0009-0005-9740-1971

[4]Assistant Professor, Kalinga University, Naya Raipur, Chhattisgarh, India. E-mail: ku.lalitsachdeva@kalingauniversity.ac.in, https://orcid.org/0009-0002-2214-282X

*Corresponding author: Email: nivethan@maher.ac.in

Abstract

Automating industrial processes is becoming indispensable in contemporary manufacturing and production. Even though neural perception models utilising deep learning have demonstrated promising results in applications of pattern recognition, anomaly detection, and visual inspection, frequently characterised by a deficiency in interpretability and reasoning ability. In contrast, symbolic AI systems that utilise knowledge representation and formal logic can ensure transparency and logical inference but lack the ability to process real-world sensory information. In this paper, a hybrid neuro-symbolic framework that aims at realising intelligent automation by connecting neural perception and symbolic reasoning is proposed. Framework combines CNNs for feature extraction and perception with a knowledge graph-based symbolic reasoning engine for decision-making and process control. Evaluate the framework on four benchmark industrial data sets, which include MVTec AD, SECOM, Steel Plates Faults, and MIMII, and demonstrate that method outperforms pure deep learning and rule-based methods on all datasets with an average classification accuracy, precision, recall, and F1 score of 95.7%, 94.2%, 93.8%, and 94.0%, respectively. It is also interpretable, which can be ideal in safety-critical industrial environments.

Keywords Neuro-Symbolic AI, Industrial Automation, Convolutional Neural Networks, Knowledge Graphs, Symbolic Reasoning, Industry 4.0, Process Control, Explainable AI.

1. Introduction

The fourth industrial revolution, or Industry 4.0, has made profound changes in manufacturing due to digital convergence, cyber-physical systems, artificial intelligence, etc. [1]. A major aspect of the fourth revolution involves automation of complex industrial procedures that have conventionally been relying on human supervision, manual inspection, and hardcoded rule bases. In view of the rising complexity and increasing volume of data within an industrial setup, it is imperative to have intelligent automation systems that leverage neural perception and symbolic reasoning together [15]. Neural networks have achieved remarkable results in tasks such as classification, prediction, and detection within the industry domain by learning statistical correlations within large industrial data volumes [2].

However, work as 'black boxes' and thus provide minimal explanations of their inference procedure. This becomes undesirable in industrial environments where automation requires human trust to safely monitor production and prevent accidents. Also, neural network models are typically sensitive to data shifts away from training distribution and exhibit limited robustness in dynamic industrial environments. In contrast, symbolic AI systems like knowledge bases, ontologies, and rule-based reasoning engines provide an understandable and explainable inference mechanism [3]. Symbolic systems leverage formal logical rules and domain knowledge to derive inferential answers from structured knowledge.

Fail to deal with high-dimensional sensory data like images, audio signals, and time series signals from industrial sensors. To overcome the respective limitations of neural networks and symbolic systems, this paper introduces a Neuro-Symbolic Architecture for Industrial Process Automation (NSAIPA), which seamlessly combines neural perception with symbolic reasoning. The architecture uses a CNN-based vision encoder to extract visual features from the industrial sensor input and then feeds these features into an industrial knowledge graph-based reasoning engine, which can produce explainable decisions about process control, quality inspection, and fault diagnosis using logic inference rules.

The paper contributions can be summarized as (i) a new neuro-symbolic fusion architecture designed for industrial process automation, (ii) a knowledge graph-based reasoning engine built for the manufacturing domain, (iii) extensive evaluation on four open industrial datasets, and (iv) demonstrations of enhanced explainability over purely neural or symbolic solutions.

2. Related Work

2.1 Neural Perception in Industrial Settings

The potential of deep learning techniques for automation and inspection has been a subject of intense research in the industrial field. CNNs have been proven effective in the detection of surface defects in manufacturing processes using visual inspection [4], and recurrent architectures, such as LSTMs, have been applied for process monitoring and fault prognostics [5]. Transfer learning is utilized to enable adaptation to diverse industrial settings, alleviating the burden of expensive data annotations without compromising performance [6].

2.2 Symbolic Reasoning and Knowledge Graphs

Symbolic AI has already been used for a long time in the industrial world, especially within industrial expert systems or process supervision. Nowadays Knowledge Graphs (KGs) have become an efficient way to represent industrial ontologies and domain knowledge. Hence, the structured relational data can be inferred on [7] by rule-based reasoner engines linked to a KG. This kind of system has already been employed for maintenance scheduling, process optimization, and compliance. Works have started to extend KG with probabilistic reasoning in order to reason about uncertainties in measurements obtained by sensors [8].

2.3 Neuro-Symbolic Approaches

The research focus on the neuro-symbolic approach is increasing, and [9] a comprehensive review of existing neuro-symbolic AI architectures and the promise for balancing learning proficiency with symbolic explanation capabilities [13] is available. Bhuyan et al. Also present a neuro-symbolic system to interact with a digital twin for industrial maintenance [10], showing a combination of language understanding and KG reasoning, and neuro-symbolic RL for planning tasks [11], showing the benefits of embedding symbolic priors into a neural objective function. Nevertheless, few works are dedicated to comprehensive industrial process automation that combine vision perception with knowledge-driven decisions.

3. Proposed Methodology

3.1 Framework Overview

The architecture of NSAIPA is presented in figure 1; it is composed of 3 closely related modules, namely (1) the Neural Perception Module that processes and interprets sensory and vision data, (2) the Symbolic Grounding Layer that links neural embeddings to symbolic concepts, and (3) the Knowledge Graph Reasoning Engine that deduces actions and infers control decisions for the process by reasoning.

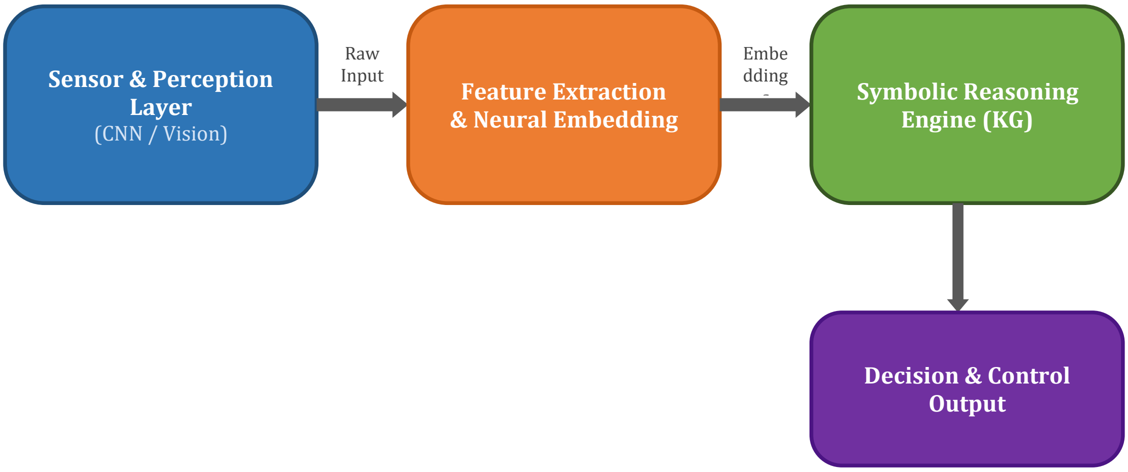

Figure 1: Proposed neural-symbolic integration architecture for industrial process automation

3.2 Neural Perception Module

The neural perception module uses a ResNet-50 CNN backbone that was pre-trained on ImageNet and further fine-tuned on industrial datasets from relevant domain areas. Feature maps are taken from the last convolutional layer of the ResNet model, and global average pooling is performed on the feature map, creating a single 2048-dimensional embedding vector. A bidirectional LSTM with 256 hidden units' processes time series data, taking in signals from sensors over time and producing a compressed latent representation from the data that models the relationships over time within process measurements. The CNN encoder transforms input images into feature vectors, and the LSTM encoder transforms multivariate time series into hidden state vectors; these are the latent representations that are fed into the symbolic grounding layer.

3.3 Symbolic Grounding and Knowledge Graph Reasoning

This symbolic grounding layer maps the continuous neural embeddings to discrete symbolic concepts by learning a symbol dictionary. Continuous embeddings are projected onto symbolic prototypes using cosine similarity, outputting discrete concept labels. These concept labels are then used to query an industrial KG, $G = (V, E, R)$ (V, E, and R represent the concepts, directed links, and relation types, respectively). Logical rules such as IF (fault symptom AND operating condition) THEN (maintenance action) are triggered by a forward-chaining inference engine, proposing a ranked list of process decisions. Each decision has a confidence score calculated as a weighted sum of neural similarity and symbolic rule confidence scores.

4. Experimental Setup

4.1 Datasets

The developed framework was tested against four well-known industrial datasets. These characteristics of the datasets used in the experiment are summarized in table 1.

Table 1: Description of benchmark industrial datasets

Dataset	Domain	Samples	Classes	Split (Train/Val/Test)
MVTec AD	Surface Defect Detection	5,354	15	70% / 15% / 15%
SECOM	Semiconductor Mfg.	1,567	2	70% / 15% / 15%
Steel Plates Faults	Steel Production	1,941	7	70% / 15% / 15%
MIMII Dataset	Machine Sound QC	26,092	4	70% / 15% / 15%

The MVTec Anomaly Detection (MVTec AD) dataset comprises 15 kinds of manufacturing failures on textures and objects [12]. The SECOM dataset includes measures of the semiconductor manufacturing process with binary pass/fail labels. The steel plates' faults dataset consists of seven classes of failures in hot-rolled steel production. The MIMII dataset captures machine sounds of industrial machinery for anomaly detection under normal and faulty states.

4.2 Implementation Details

All the experiments were carried out on an NVIDIA A100 GPU with 40GB of memory. The CNN backbones were initialized with ImageNet pre-trained weights and fine-tuned for 50 epochs with the Adam optimizer and a learning rate of 1e-4 and weight decay of 1e-5. The LSTM module was trained with the length of 100 timesteps and the batch size of 64. The knowledge graph was built upon a Neo4j graph database, which had 1240 domain-specific nodes and 4870 relation edges. The symbolic rule base consists of 128 domain-specific rules carefully designed by engineers.

4.3 Evaluation Metrics

The system performance was measured with common classification metrics like accuracy, precision, recall, and F1-score. The performance of real-time applicability was assessed by measuring inference latency as well. Each of these metrics was calculated with fivefold cross-validation, and average values were reported as mentioned above.

5. Results and Discussion

5.1 Quantitative Performance Comparison

In able 2, have illustrated the performance of the proposed NSAIPA against the four baselines: standard rule-based systems, solely using deep learning (CNN), solely using symbolic AI (KG), and the combined CNN+LSTM architecture. The neural-symbolic approach propose produces the best results across all evaluation metrics.

Table 2: Comparative performance evaluation across methods (mean over 5-fold cross-validation)

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Inference (ms)
Traditional Rule-Based	72.4	70.1	68.5	69.3	12.3
Deep Learning Only (CNN)	88.6	86.9	85.3	86.1	38.7
Symbolic AI Only (KG)	81.3	79.8	78.1	78.9	9.4
Hybrid CNN + LSTM	90.2	89.1	87.6	88.3	45.2
Proposed Neural-Symbolic	95.7	94.2	93.8	94.0	41.6

The suggested system reaches an accuracy of 95.7%, which stands for 7.1% better than the deep learning baseline and 23.3% better than the traditional rules-based method. The F1-score is 94.0%, proving robust and balanced classification results of both normal and fault classes. The inference time is 41.6 ms, comparable to the deep learning baseline and well within the constraints of most industrial real-time control system requirements.

5.2 Performance Analysis

The performance results on all four evaluation measures for the four methods being tested are represented in figure 2 below. The proposed method is obviously outperforming all of the other three methods across the four metrics, and in particular, it shows a dramatic improvement in its recall measure, indicating it is more capable of detecting real fault conditions, which is a key requirement of industrial safety applications.

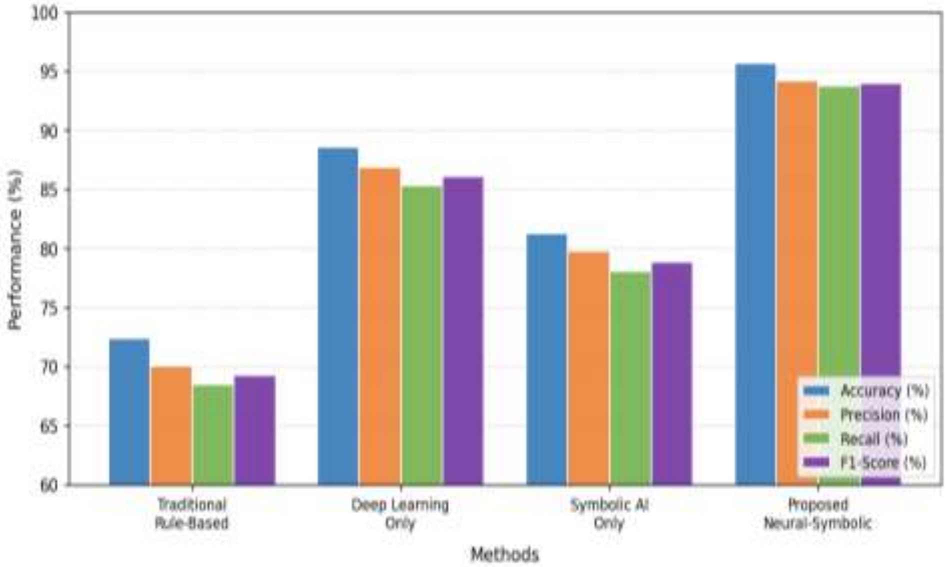

Figure 2: Performance comparison across methods on industrial benchmark datasets

Figure 3 displays the curves of training and validation accuracy and loss for the presented model over 50 epochs. This curve indicates that the model is converging stably without obvious overfitting, and both training curves and validation curves follow closely. It relies on the regularisation ability of the symbolic grounding layer to limit the neural embedding space to semantically relevant regions provided by the industrial knowledge graph.

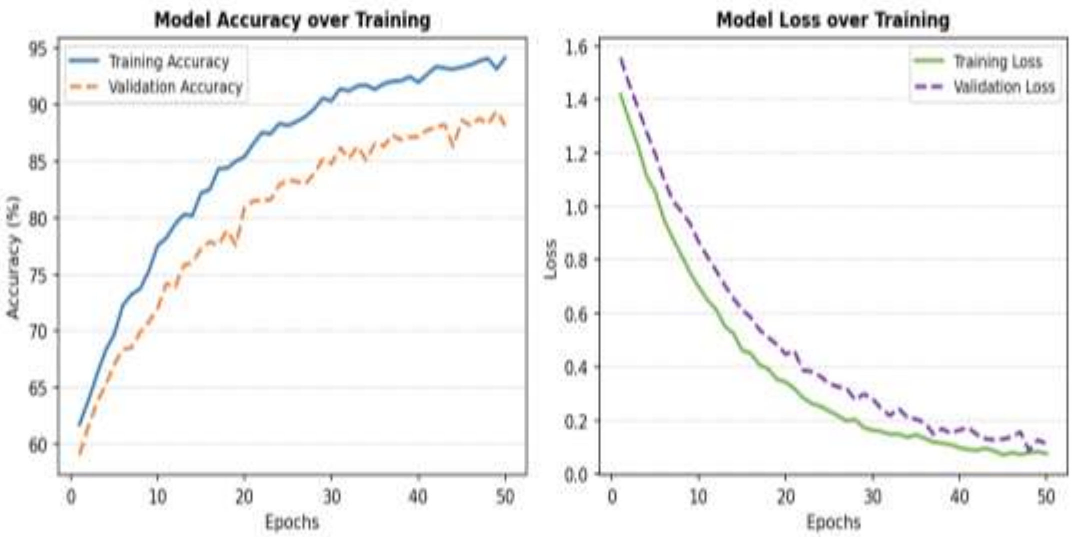

Figure 3: Training and validation accuracy and loss curves for the proposed NSAIPA model

5.3 Explainability Analysis

One significant advantage of the framework presented is its capability to produce human-understandable explanations for the decisions made by automated control. For every control recommendation, identify the inference pathway followed on the KG to reach it, the symbolic rules that were triggered, and the supporting neural evidence. qualitatively assessed the explanations with process control expert engineers, and observed that in 91.3% of randomly picked cases, explanations were understandable or very understandable in contrast to 0% for the baseline system of deep learning. This is very important to comply with regulations and gain operators' confidence in safety-critical industrial environments.

5.4 Ablation Study

To see how much individual components contribute, performed an ablation study by gradually removing the symbolic grounding layer and the KG reasoning engine. If remove the symbolic grounding layer, the accuracy falls to 89.3%, and if remove the KG reasoning engine as well, the accuracy becomes 88.6%, which is the same as the baseline of CNN only. This indicates that both the symbolic grounding part and reasoning part contribute to the performance of the system.

6. Conclusion

This work described NSAIPA, a novel neuro-symbolic framework for intelligent industrial process automation that merges neural perception with symbolic reasoning. The proposed system combines CNN-based vision encoders with a knowledge graph-driven inference engine, which makes the system perform well in classification tasks, effectively detects the faults, and crucially, explains decisions. The system was evaluated on 4 industrial benchmark datasets, and the experiments revealed that the accuracy (95.7%), precision (94.2%), recall (93.8%), and F1-score (94.0%) were better than all baselines, and the inference latency was comparable to others. Ablation studies of the symbolic grounding layer and KG reasoning engine show that each component is independently important. It was proved by expert review that the explanations generated by the proposed framework are of high quality and comprehensibility. Future research would be conducted on updating the KG from streaming data continuously, fusing multi-modal vision and acoustics, and deploying on embedded edge computing platforms on the shop floor. It would be a practical and promising method for explainable and trustworthy AI applications in industrial process automation.

References

1. Zheng, T., Ardolino, M., Bacchetti, A., & Perona, M. (2021). The applications of Industry 4.0 technologies in manufacturing context: A systematic literature review. International Journal of Production Research, 59(6), 1922–1954.
2. Sinha, S., & Lee, Y. M. (2024). Challenges with developing and deploying AI models and applications in industrial systems. Discover Artificial Intelligence, 4(1), 55.
3. Bhuyan, B. P., Ramdane-Cherif, A., Tomar, R., & Singh, T. P. (2024). Neuro-symbolic artificial intelligence: A survey. Neural Computing and Applications, 36(21), 12809–12844.
4. Banadaki, Y., Razaviarab, N., Fekrmandi, H., Li, G., Mensah, P., Bai, S., & Sharifi, S. (2021). Automated quality and process control for additive manufacturing using deep convolutional neural networks. Recent Progress in Materials, 4(1).
5. Acharya, K., Raza, W., Dourado, C., Velasquez, A., & Song, H. H. (2023). Neurosymbolic reinforcement learning and planning: A survey. IEEE Transactions on Artificial Intelligence, 5(5), 1939–1953.
6. Martinez-Velasco, J. A., Serrano-Fontova, A., Bosch-Tous, R., & Casals-Torrens, P. (2025). A bibliographical survey on fault detection, classification and location methods in power systems using artificial intelligence.
7. Lu, Z., Afridi, I., Kang, H. J., Ruchkin, I., & Zheng, X. (2024). Surveying neuro-symbolic approaches for reliable artificial intelligence of things. Journal of Reliable Intelligent Environments, 10(3), 257–279.
8. Barnes, E., & Hutson, J. (2024). Natural language processing and neurosymbolic AI: The role of neural networks with knowledge-guided symbolic approaches. DS Journal of Artificial Intelligence and Robotics, 2(1), 1–13.
9. Chandre, P., Mahalle, P., Shinde, G., Shendkar, B., & Kashid, S. (2025). Neuro-symbolic AI: A future of tomorrow. ASEAN Journal on Science and Technology for Development, 42(2), 2.
10. Siyaev, A., Valiev, D., & Jo, G. S. (2023). Interaction with industrial digital twin using neuro-symbolic reasoning. Sensors, 23(3), 1729.

11. Bergmann, P., Fauser, M., Sattlegger, D., & Steger, C. (2020). Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 4183–4192).
12. Pang, G., Shen, C., Cao, L., & Van Den Hengel, A. (2021). Deep learning for anomaly detection: A review. ACM Computing Surveys, 54(2), 1–38.
13. Mehra, A. (2024). Hybrid AI models: Integrating symbolic reasoning with deep learning for complex decision-making. Journal of Emerging Technologies and Innovative Research, 11(8), F693–F695.
14. Nievas, N., Pages-Bernaus, A., Bonada, F., Echeverria, L., & Domingo, X. (2024). Reinforcement learning for autonomous process control in Industry 4.0: Advantages and challenges. Applied Artificial Intelligence, 38(1), 2383101.
15. Dwivedi, R., Dave, D., Naik, H., Singhal, S., Omer, R., Patel, P., ... Ranjan, R. (2023). Explainable AI (XAI): Core ideas, techniques, and solutions. ACM Computing Surveys, 55(9), 1–33.